arpa·e
CHANGING WHAT'S POSSIBLE

Energy Efficiency Computing Workshop
"Cooling Compute Systems Efficiently, Anytime, Anywhere"

Peter de Bock, Program Director
Vivien Lecoustre, Tech SETA
Carlos Noyes, Tech SETA
Tom Bress, Tech SETA
Rakesh Radhakrishnan, T2M Advisor
John Gallagher, T2M Summer Scholar

Datacenters energy consumption

- The **United States dominates the Data Center industry** and are **poised to grow (CAGR of 20%)** both in the number of centers and as a portion of the United States' energy consumption
- Data Centers currently **consume between 2-5% of the entirety of the US electricity** (this represents between 0.3 to 0.7 quads of electricity; 0.9 to 2.1 quads of primary energy)

Number of data centers worldwide in 2021, by country

Number of data centers worldwide 2021, by country

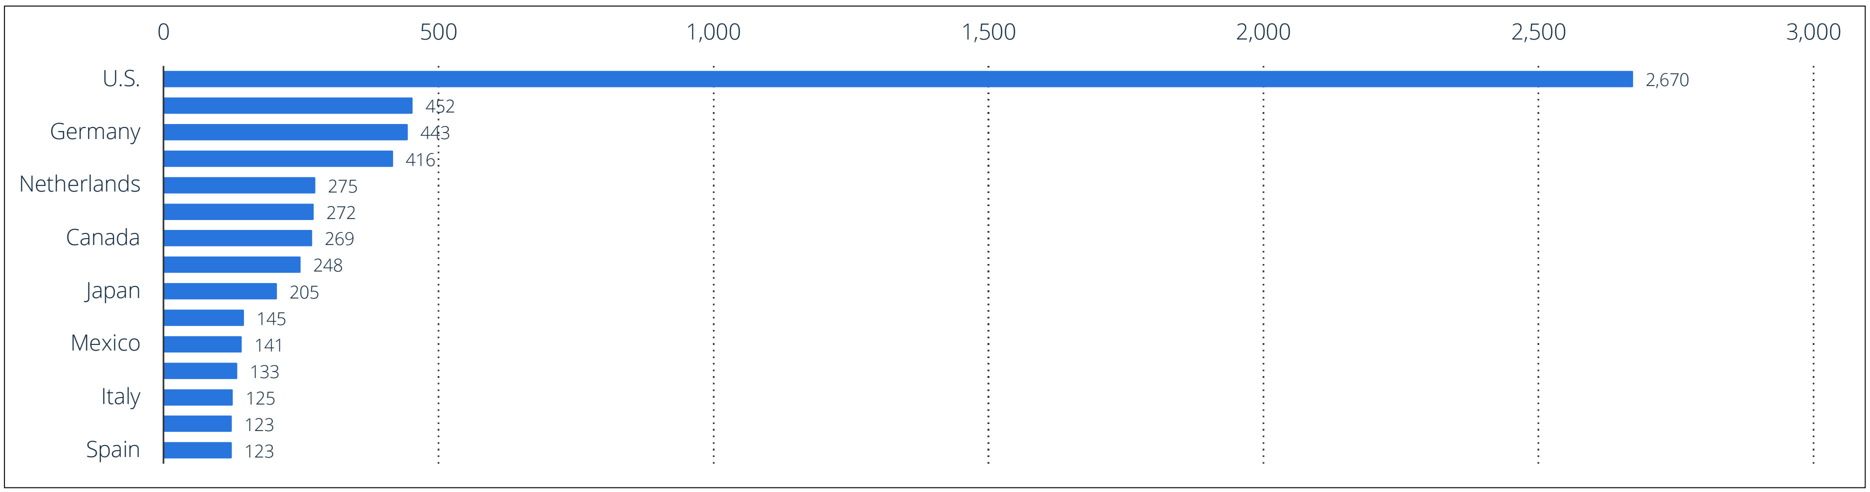

Source(s): Cloudscene; ID 1228433

CHANGING WHAT'S POSSIBLE

Baseline projections for 2030

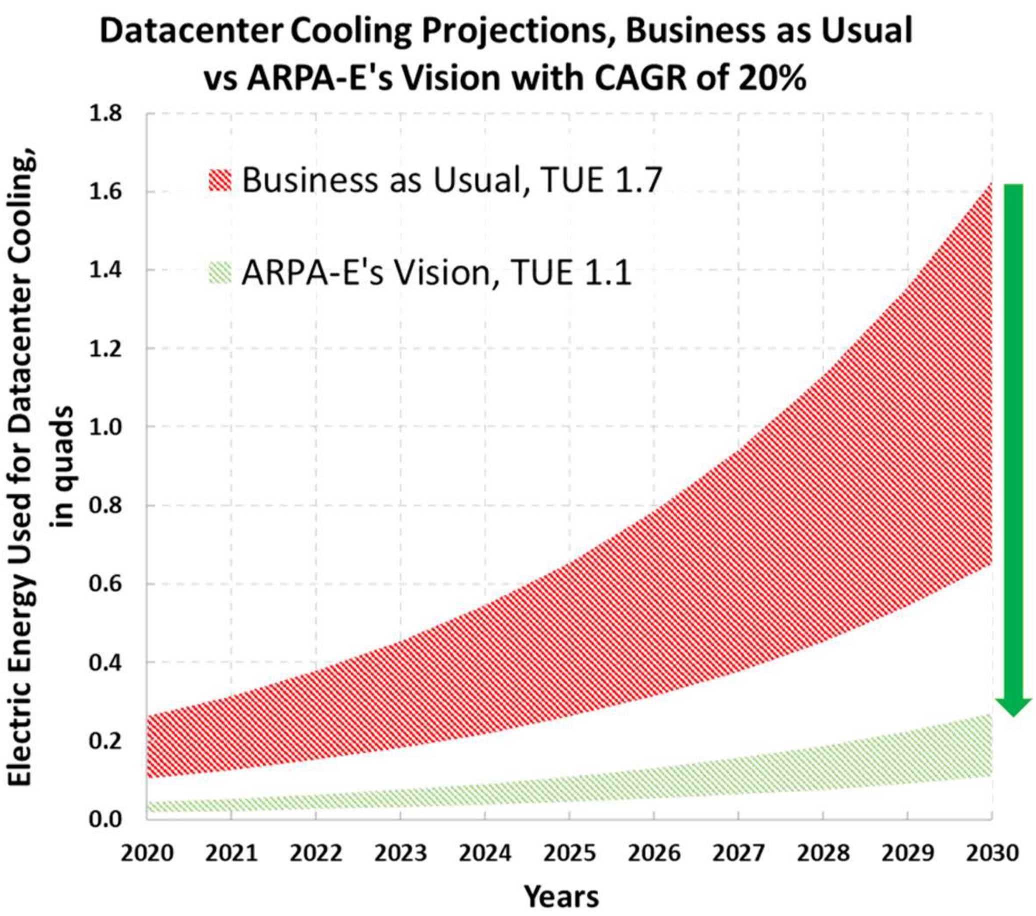

Although projections are diverse, this is a growing sector

Potential to **save between 0.5 to 1.4 quads of electricity** by **transformative improvements** of datacenter cooling by 2030.

This represents **between 1.6 to 4.1 quads of primary energy** saved[1].

We do not wish to use Energy for cooling – it does not add value

[1]*Assuming 35% of generation efficiency & 95% transmission efficiency*

Electronics cooling

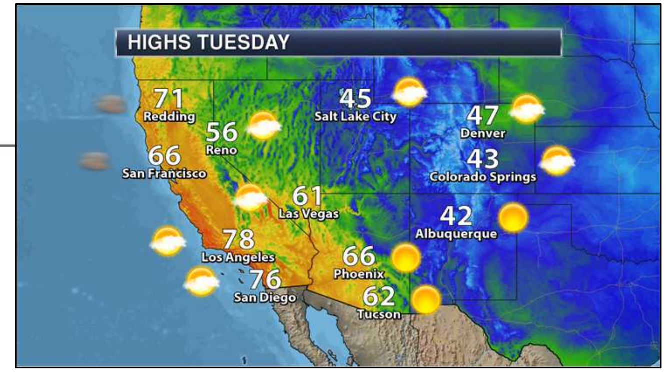

Electronics run at **70-90°C (160-194°F)** reliably

> *"2nd law of Thermodynamics: Heat always flows spontaneously from a hotter to a colder body"*

Outside **ambient temperatures** are **<u>much lower</u>** than **electronics temperatures**; so why do we use massive amounts of **Energy and Water** to reject heat to ambient?

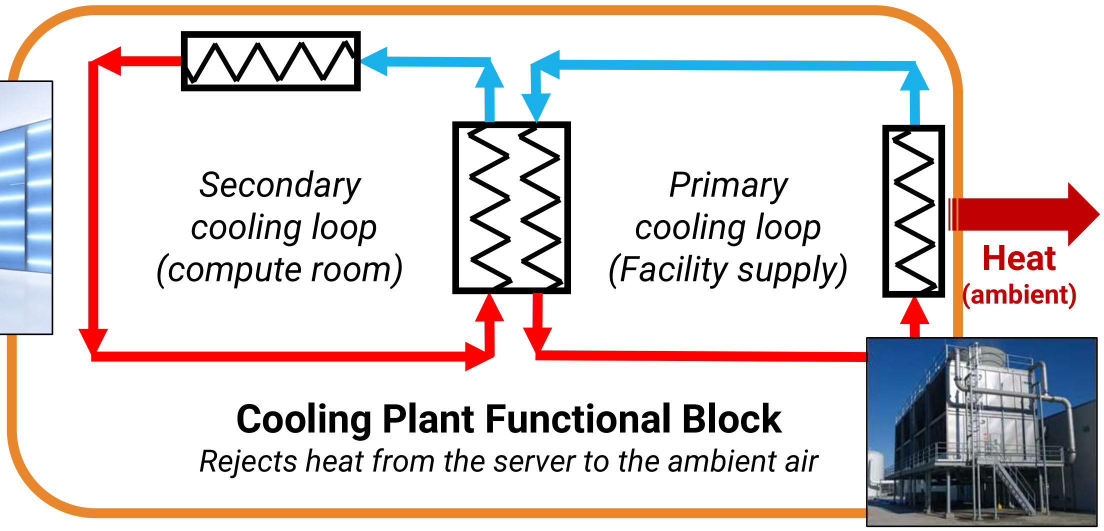

- Chiller
- Evap. Cooler
- Dry Cooler

Heat rejection from data centers

Inefficiencies and **rising power densities** in the **compute room cooling loop** ….

…. Force the **facility supply temperature** to be **set to lower temperatures**

Thermal Resistance = Temperature rise per Watt rejected [K/W]

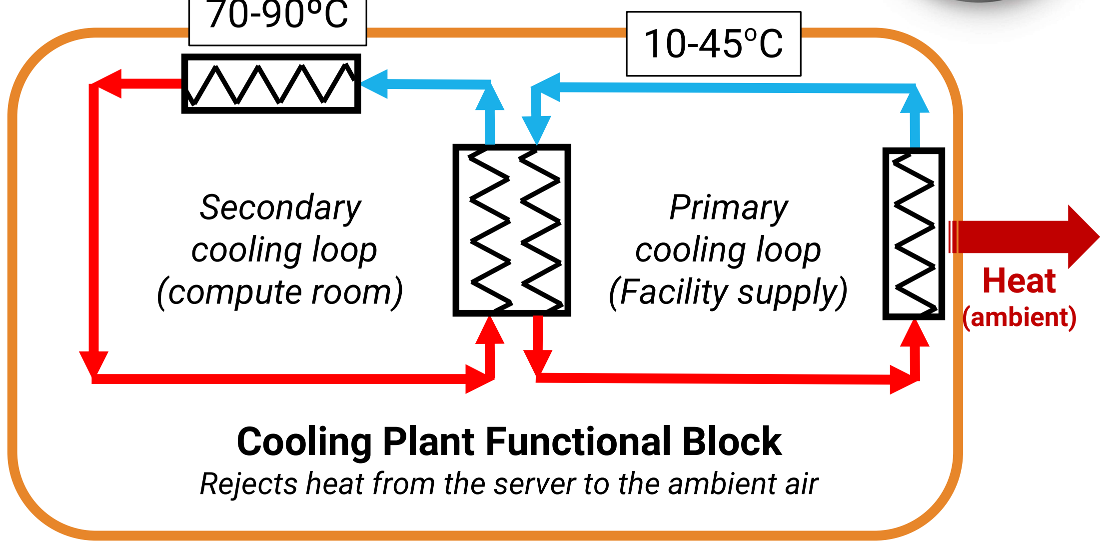

The lower the thermostat, the more energy we use ☹

Chiller

Evap. Cooler

Dry Cooler

We're out of luck, chip and server '25-'30 power projected up sharply

With Moore's law sunsetting the benefits of **transistor scaling** are **diminishing**. Chip architectures of the future are **diverse** and **higher power**. **AI & ML** add a diversity of onboard GPU/XPU cores using **additional power**

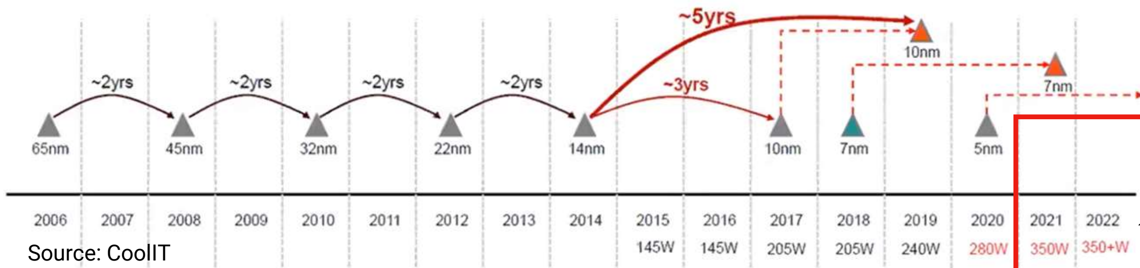

Source: CoolIT

2025-2030
1000W/chip?
+ 8-16 GPUs?
6kW/server?

High Power Chip Era "★Wild West★" $T_j\downarrow$

Nvidia A100 GPU 400W (server can have 8x400W in addition to main CPU)

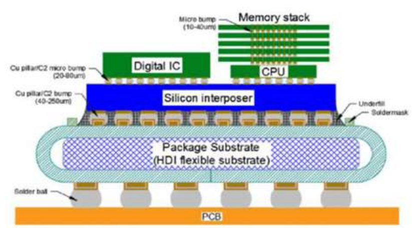

Heterogenous Integration stacking chiplets in 3D configuration challenging to cool!

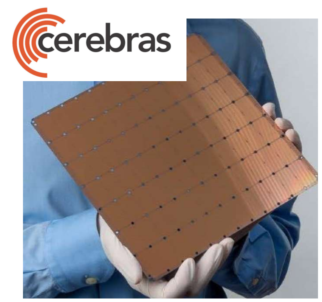

Analog computing Cerebras 71 in^2

rise in FPGAs projected optimized for application

Cooling difficulty projected to rise

Emergence and Expansion of Liquid Cooling in Mainstream Data Centers, ASHRAE White Paper, Technical Committee 9.9, **2021**

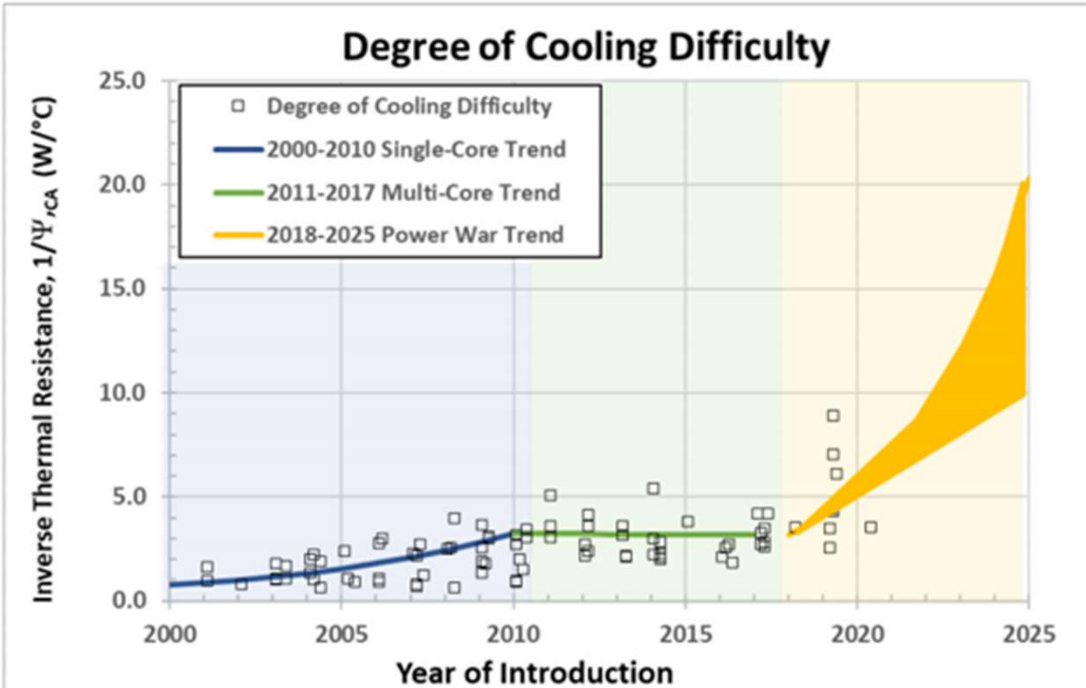

ASHRAE coming out with W17, W27, W32, W40, W45, W+ standards

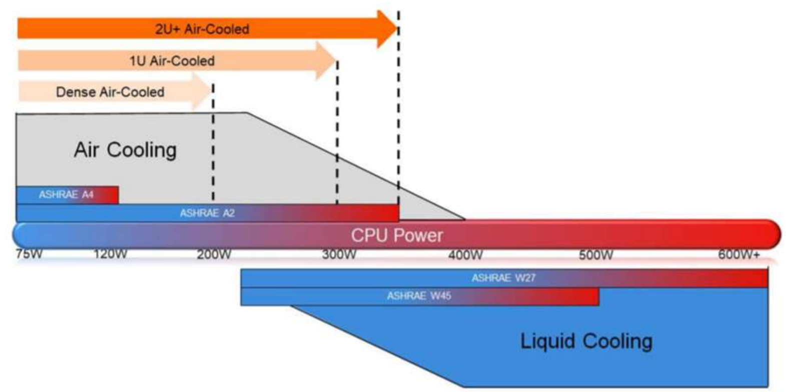

… but predicts that high power CPUs (over 500W) should use W27(<27℃) or less Facility supply temperature

… as server power goes up, facility supply temperatures projected to need to go down due to cooling difficulty, = more energy for cooling

Cooling difficulty projected to rise

Emergence and Expansion of Liquid Cooling in Mainstream Data Centers, ASHRAE White Paper, Technical Committee 9.9, **2021**

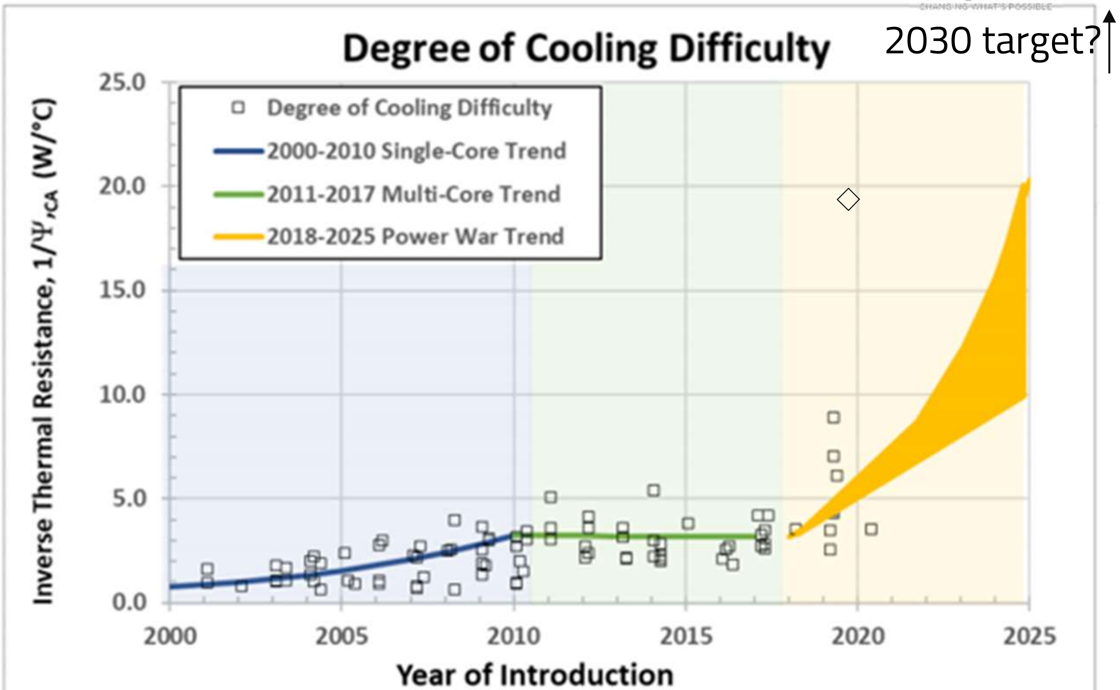

Current chart depicts conductance from chip to coolant in secondary loop.
Potential impact target needs to reflect rejection to primary loop.

◇ **Best in-class liquid cooled thermal resistance $\theta = 0.05$ °K/W**

https://www.datacenterknowledge.com/industry-perspectives/skived-coldplates-technical-brief

Vision: If **technology is developed** to reject heat from future servers **10 x more efficiently** to facility supply, **cooling energy is saved**

Setting the supply thermostat high above ambient, closer to electronics temperature:

- AC/Chiller minimally/not needed (perhaps just for humans in room, optional)
- Hot Dry cooler >90% less power
- Hot Evap. coolers > 90% less fan power

*) Compared to 35C rejection baseline

Heat rejection from data centers

If **heat rejection** **improves greatly**, and becomes as safe, reliable and cost as **air cooling**…

…. **Facility supply** can be much warmer ambient heat rejection is **easy**

Minimal Energy use ☺

Lower Thermal Resistance 10 x

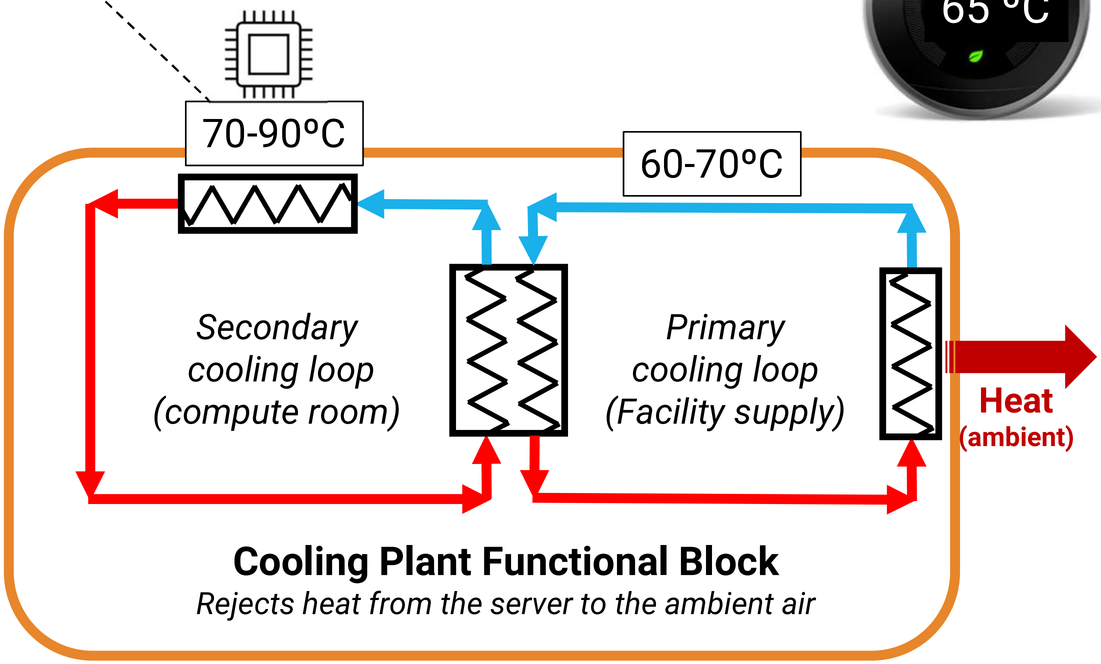

Chiller

Evap. Cooler

Dry Cooler

Facility supply temperature key driver for energy use

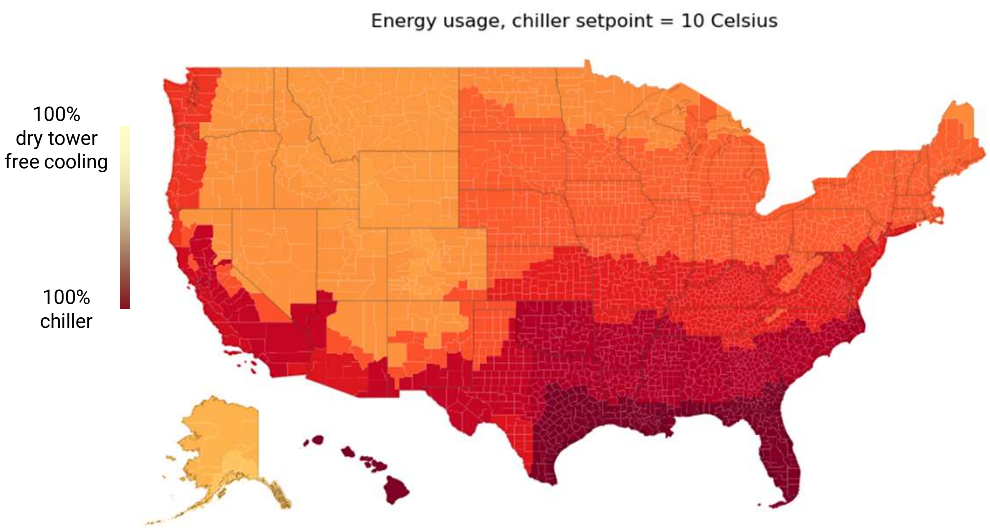

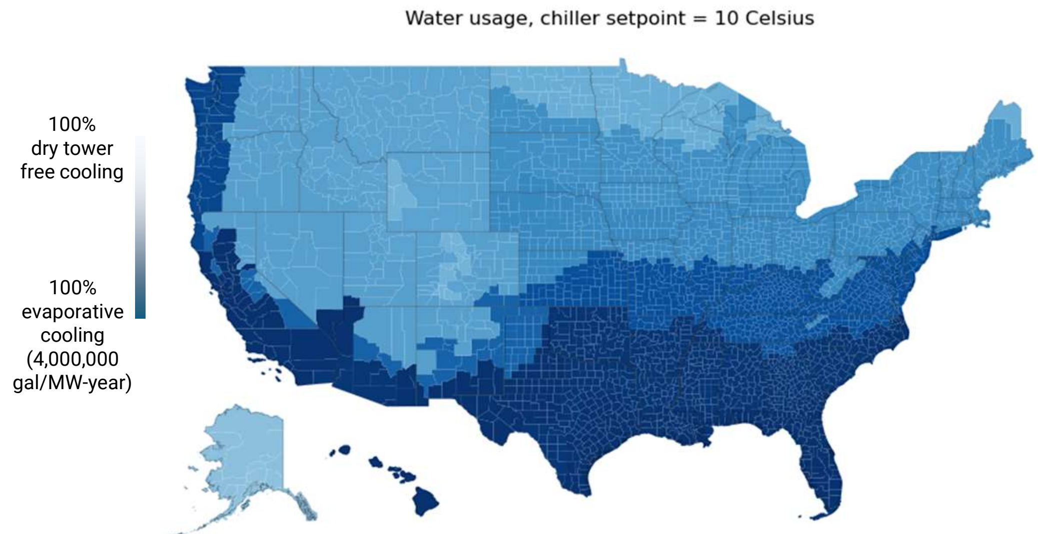

Common facility temperatures 10-32ºC. Example to illustrate effect of facility Temperature on Energy use

- The price paid for the standard supply temperature low is **Excessive Energy Use**
  - Chiller has to run most of the year – 0.75 quads for cooling
  - Water is consumed in most locations – *approx. total of >500 billion gallons of water use attributable to US data centers (~57% sourced from potable water)*

https://www.nature.com/articles/s41545-021-00101-w

Efficient heat rejection can Change the Landscape

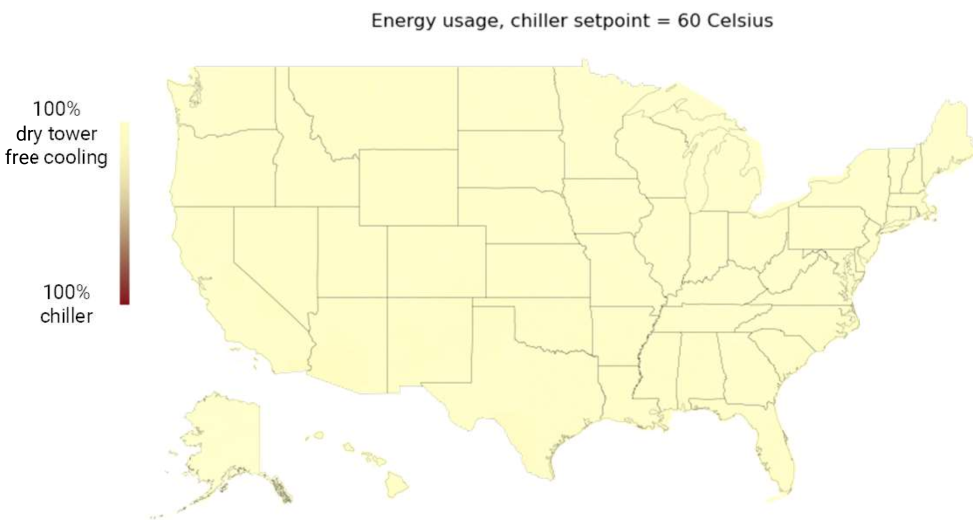

If **technology is developed** to reject heat from future servers **10 x more efficient in secondary loop (chip to facility supply)**; facility temperatures can be evaluated, and **cooling energy is saved**

<u>Bonus features:</u>

+ Location/climate independence

+ Minimal/No need for water usage

+ Reduced footprint

+ Heat rejection >60°C facilitates future WHR

How do we achieve Impact?

Innovation acceptance equation:
(derived from kano model)

Has to be leap >10x
-> performance criteria

$$\frac{\text{Performance}}{\text{Risk, Cost}} + \text{wows}$$

Some unique capabilities in the new technology
-> delighters

Must have framework: Risk and cost have to stay at parity or reduce (at system level)
-> acceptance criteria

https://en.wikipedia.org/wiki/Kano_model

"No one ever got fired because they used too much energy, people get fired when DC goes down"

Impact: we need to get to the point that by 2025-2030 the most cost-effective data center is also the most energy efficient

LCOC - Levelized Cost of Cooling (Day II focus)

$$\frac{\$}{W_{rej}}$$

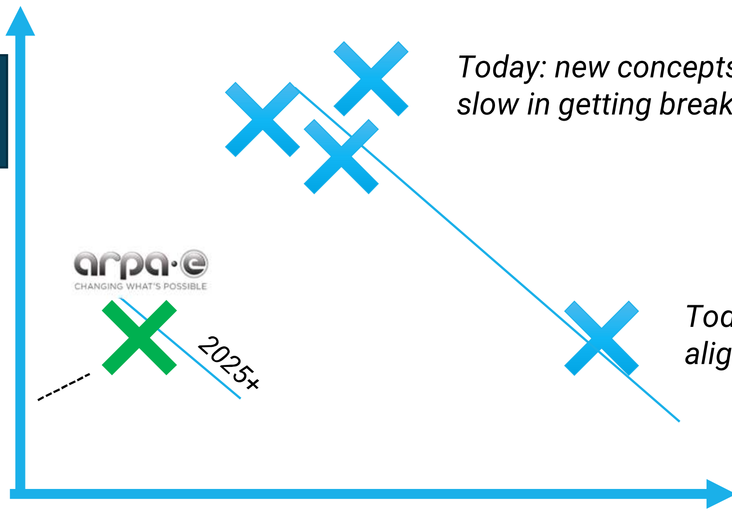

Today: new concepts for diverse reasons slow in getting breakthrough market share

Today: most cost effective does not align with most energy efficient

*New solutions need to be energy efficient **<u>and</u>** equally or better cost competitive for EDGE, co-location or larger scale*

$$\frac{W_{cooling}}{W_{rej}}$$

1/(COP) – Performance of the cooling plant (Day I focus)

Other dimensions: reliability, safety apply also

What is disruptive?! What is our moon shot?

Efficient computing vision &holy grail: (refined during this workshop)

Performance

- 1/COP: Energy of cooling less than 3%?-0%? of energy rejected
- Capable of compute systems ‘25-’30; i.e. 100-200 kW/rack @ 1000W/chip? Works on any xpu (chip, memory, gpu, etc.)
- Heat capture few < 5°C below chip temp? > 60°C exit for WHR?
- Co-designed reliable system equivalent to today's air cooling

Must haves / Bounds

- As easy to manage, maintain and operate as air-cooled data center
- Reliability similar to existing systems.
- LCOC: Cost effective → most energy efficient solution is also most cost-effective

Wows

- Location independent cooling, 24/7/365 anywhere
- Footprint of mechanical plant/BOP 1:1 to racks
- No\Minimal water usage
- Modular, pre-fab as efficient as large data centers, EDGE?
- System Autonomy? (if it allows for energy reduction)?

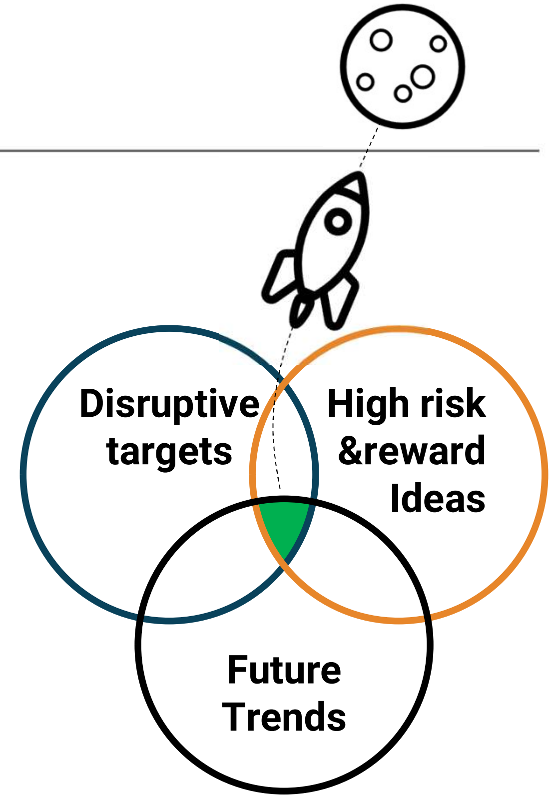

Achieving reliability will require System, Sensors& Controls Co-Design Vision

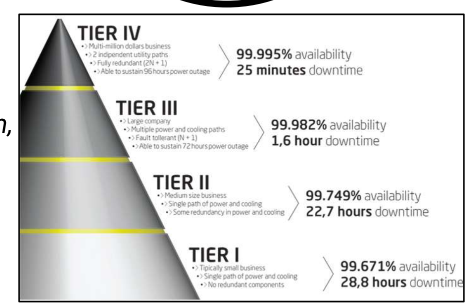

https://www.netari.com/post/2014/02/04/what-to-look-for-in-a-data-center-understanding-tier-levels-industry-standards

Potential Scope Concepts

Out of Technical Innovation scope for this effort

- Power Electronics
- Heat Generating device

Technical Innovation Scope B: From chip surface to ambient within constrained dimensions (modular)

Technical Innovation Scope A: From chip surface to facility supply conventional data hall (retrofit or new build)

+BOP: Anything that is needed for secondary ***system*** *to work* *reliably: innovation in pumps, baths, racks, fluids, connectors, hoses, filters, sensors, control co-design, logic, CDU, HX, etc.*

1. Near-device heat extraction
2. To Facility supply heat rejection

+BOP: Building shape, primary loop pumps, layout, dry coolers, etc. cannot exceed dimensions ISO 40 container

3. Heat rejection to ambient

Technical Scope A: Secondary Loop

Chip (and other devices) surface to facility supply by any means, but meeting potential program technical, operational and cost requirements

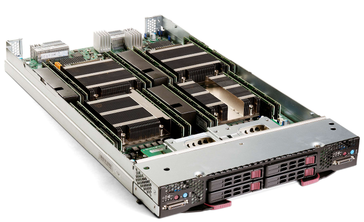

Technical Scope B: Secondary + Primary loop

System-level approach, "All-in-One" Edge data center

Will simplified cooling(Scope A) and hot rejection enable modular DC?
- Self contained
- EDGE high power density

Best in-class 100-150kW/ISO40
- ARPA-E Target 1MW in ISO40?

https://datacentrereview.com/2021/06/schneider-reveals-all-inone-liquid-cooled-ecostruxure-modular-data-centre/

Scope of Discussions: Category A & Category B

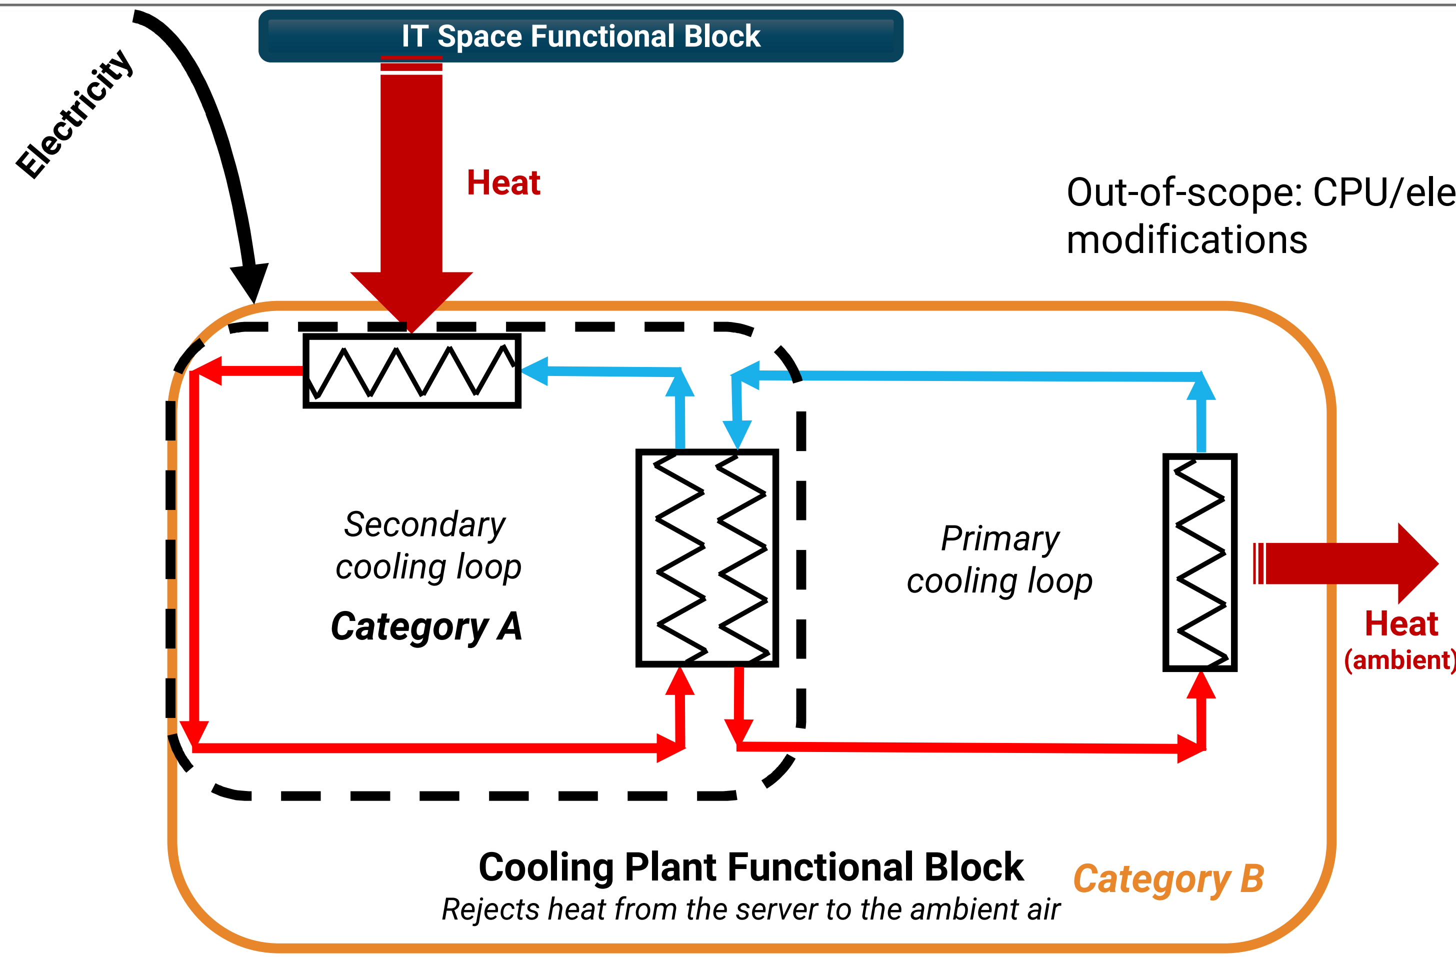

Category B: Modular data center as Tech/EDGE development platform

Tech Development Platform – System approach

- Cooling innovation (cat A) lead to location independence, less/no water need = enabler for modular systems?
- Allows for volume constrained tech development within ISO40' container form factor (320 ft^2), currently best in-class 100-150kW
- ARPA-E hard? : i.e. **system-level** approach for:
  - ~1 MW, very efficient (TUE lower than 1.1 = 91% energy used to power XPUs). **Compute density of at least 3.28 kW/sq ft**.
  - Prefabricated modular system with hot heat rejection with potential for useable or monetizable high-grade heat for WHR heating applications.
  - Vision: Self contained except for power, minimal/no water usage. Easy to install in locally.
  - Reliable: adoption path lead to Tier2? -Tier 3: 99.98%?

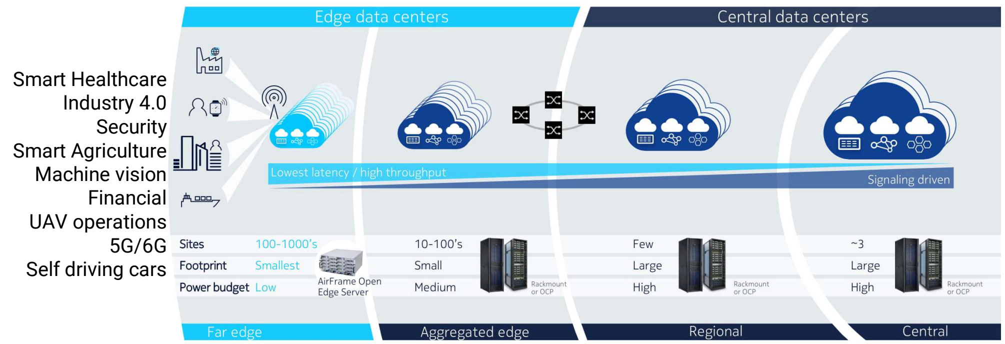

Low latency and security close to customer

https://networks.nokia.com/sites/content/files/openedge_architecture_0.jpg

Modular Data Center

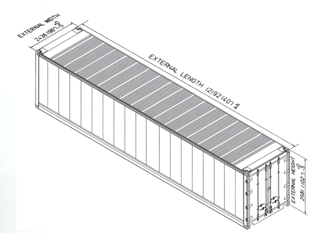

< ISO 40' container

https://datacentrereview.com/2021/06/schneider-reveals-all-inone-liquid-cooled-ecostruxure-modular-data-centre/

Volume constraint makes this an interesting tech development platform, tech could proliferate back to large data center

Let's discuss in the breakout!

How do we frame such a challenge? Breakouts

Wows

+ Location/climate independence
+ Reduced footprint
+ No need for water usage

+ Heat rejection >60°C facilitates future WHR

= Enabling modularity?

Performance

Transformational Cooling Energy reduction

How we get there (not discussed today)

Where we are today

Framework / Must-haves Bounds

Reliability

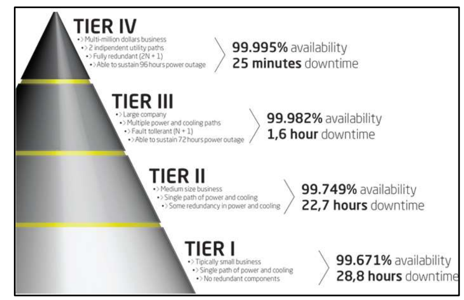

Achieving reliability will require System, Sensors& Controls Co-Design Vision

Cost

Safety

Other?

This is our opportunity to make difference

Industry choices
Rapid growth
Processor power↑

Traditional research

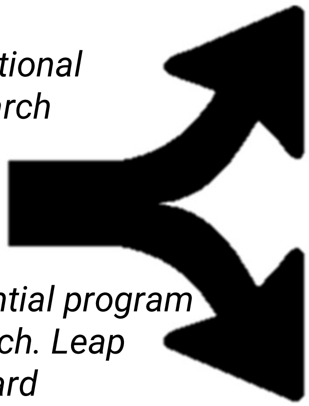

Potential program to Tech. Leap forward

Low risk: use more Energy and Water (both cheap)
Run chillers lower, add more chillers

X multiply by trends

Usurping Energy

Regional Inequality

CHANGING WHAT'S POSSIBLE

Develop leap in efficient heat rejection with similar reliability, cost and operation to air

Transformational trend breach
US leadership in Energy Efficient Computing
- Domestic market and exports

This is our opportunity to make difference